arXiv:1303.2071v1 [cs.CV] 8 Mar 2013

Application of the SP theory of intelligence to the understanding of natural vision and the development of computer vision

J Gerard Wolff*

6 March 2013

Abstract

The *SP theory of intelligence* aims to simplify and integrate concepts in computing and cognition, with information compression as a unifying theme. This article discusses how it may be applied to the understanding of natural vision and the development of computer vision. The theory, which is described quite fully elsewhere, is described here in outline but with enough detail to ensure that the rest of the article makes sense.

Low level perceptual features such as edges or corners may be identified by the extraction of redundancy in uniform areas in a manner that is comparable with the run-length encoding technique for information compression.

The concept of *multiple alignment* in the SP theory may be applied to the recognition of objects, and to scene analysis, with a hierarchy of parts and sub-parts, and at multiple levels of abstraction.

The theory has potential for the unsupervised learning of visual objects and classes of objects, and suggests how coherent concepts may be derived from fragments.

As in natural vision, both recognition and learning in the SP system is robust in the face of errors of omission, commission and substitution.

The theory suggests how, via vision, we may piece together a knowledge of the three-dimensional structure of objects and of our environment, it provides an account of how we may see things that

*Dr Gerry Wolff, BA (Cantab), PhD (Wales), CEng, MBCS (CITP); CognitionResearch.org, Menai Bridge, UK; jgw@cognitionresearch.org; +44 (0) 1248 712962; +44 (0) 7746 290775; www.cognitionresearch.org.

are not objectively present in an image, and how we recognise something despite variations in the size of its retinal image. And it has things to say about the phenomena of lightness constancy and colour constancy, the role of context in recognition, and ambiguities in visual perception.

A strength of the SP theory is that it provides for the integration of vision with other sensory modalities and with other aspects of intelligence.

Keywords: vision, information compression, artificial intelligence, perception, cognition, representation of knowledge, learning, pattern recognition, natural language processing, reasoning, planning, problem solving.

1 Introduction

The *SP theory of intelligence* aims to simplify and integrate ideas in artificial intelligence, mainstream computing, and human perception and cognition, with information compression as a unifying theme. The theory is described in several peer-reviewed articles,[1] and most fully in Wolff (2006).

The main purpose of this article is to describe how the SP theory may be applied to the understanding of natural vision and the development of computer vision, and to discuss associated issues. Both of those themes—natural vision and artificial vision—are discussed together throughout the article, since each one may illuminate the other.

In broad terms, the potential benefits of the SP theory in those two areas are the simplification and integration of concepts, deeper insights, better performance (of artificial systems), and the seamless integration of vision with other sensory modalities, and with other aspects of intelligence such as reasoning, planning, problem solving, and unsupervised learning. What is perhaps the main attraction of the theory is the potential for one relatively simple framework to accommodate several different aspects of intelligence, including vision.

As a preliminary, the next section describes the theory in outline, with associated ideas.

2 Outline of the SP theory

The SP theory combines conceptual simplicity with descriptive and explanatory power in several areas, including concepts of 'computing', the repre-

[1] See www.cognitionresearch.org/sp.htm and links from there.

sentation of knowledge, natural language processing, pattern recognition, several kinds of reasoning, the storage and retrieval of information, planning and problem solving, unsupervised learning, information compression, and human perception and cognition.

Since the SP theory has been described quite fully in (Wolff, 2006), only the essentials will be given here, with enough detail to ensure that the rest of the article makes sense.

The main elements of the SP theory are:

- The theory is conceived as an abstract system that, like a brain, may receive 'New' information via its senses and store some or all of it as 'Old' information.
- All New and Old information is expressed as arrays of atomic symbols (*patterns*) in one or two dimensions.
- The system is designed for the unsupervised learning of Old patterns by compression of New patterns.
- An important part of this process is, where possible, the economical encoding of New patterns in terms of Old patterns. This may be seen to achieve such things as pattern recognition, parsing or understanding of natural language, or other kinds of interpretation of incoming information in terms of stored knowledge, including several kinds of reasoning.
- Compression of information is achieved via the matching and unification (merging) of patterns, with key roles for the frequency of occurrence of patterns, and their sizes.
- The concept of *multiple alignment*, outlined in Section 2.2, is a powerful central idea, similar to the concept of multiple alignment in bioinformatics but with important differences.[2]
- Owing to the intimate connection between information compression and concepts of prediction and probability (see, for example, Li and Vitányi, 2009), it is relatively straightforward for the SP system to calculate probabilities for inferences made by the system, and probabilities for parsings, recognition of patterns, and so on.

[2]For a summary of the differences, see Wolff (2006, Section 3.4.1).

- In developing the theory, I have tried to take advantage of what is known about the psychological and neurophysiological aspects of human perception and cognition, and to ensure that the theory is compatible with such knowledge. The way the SP concepts may be realised with neurons (*SP-neural*) is discussed in Wolff (2006, Chapter 11).

2.1 Computer models

The SP theory is realised in the form of computer models which may be regarded as first versions of the *SP machine*, an expression of the theory and a means for it to be applied. The SP70 model is the most comprehensive version, with capabilities in the building of multiple alignments and unsupervised learning. The SP62 model is the same but it lacks any ability to learn. Although SP62 is a subset of SP70, it has proved convenient to maintain them as separate models.[3]

At the heart of the SP models is a process for finding good full or partial matches between patterns (Wolff, 2006, Appendix A), with a flexibility that is somewhat like the WinMerge utility for finding similarities and differences between files, or standard 'dynamic programming' methods for the alignment of sequences. The main difference between the SP process and others, is that the former can deliver several alternative matches between patterns, while WinMerge and standard methods deliver one 'best' result.

Multiple alignments are built in stages, with pairwise matching and merging of patterns, and with merged patterns from any stage being carried forward to later stages. At all stages, the aim is to encode New information economically in terms of Old information and to weed out multiple alignments that score poorly in that regard.

In the SP70 model, there are additional processes for deriving Old patterns from multiple alignments, evaluating sets of newly-created Old patterns in terms of their effectiveness for the economical encoding of the New information, and weeding out low-scoring sets.

More detail about SP70 may be found in Wolff (2006, Sections 3.9 and 9.2). The SP61 model, a precursor of SP62 which is very similar to it, is described in Sections 3.9 and 3.10 (*ibid.*).

The main limitations of current models are:

- That they work with one-dimensional patterns and have not yet been generalised to work with 2D patterns (although a preliminary attempt

[3]The source code for the SP62 and SP70 computer models, with associated documents and files, may be downloaded via links under the heading 'SOURCE CODE' at the bottom of the page on bit.ly/WtXa3g.

has been made to consider how the SP principles may be generalised to patterns in two dimensions (Wolff, 2006, Section 13.2.1)).

- That the arithmetic meaning of numbers is not recognised—they are simply treated as patterns.
- That SP70 does not yet learn intermediate levels of abstraction in grammars, or discontinuous patterns in data.

I believe these problems are soluble. Potential solutions will be mentioned at relevant points below. Owing to the first of these limitations, most of the examples in this article, and much of the discussion, will relate to one-dimensional patterns.

2.1.1 Computational complexity

Like most problems in artificial intelligence, the problems that are addressed in the SP models—finding good full and partial matches between patterns, the formation of multiple alignments, and the learning of useful sets of patterns—are not tractable[4] if the requirement is to find ideal solutions. But, as with most programs in artificial intelligence, things become much easier if one is content with solutions that are reasonably good and not necessarily perfect.

Like most programs in artificial intellegence, the SP models apply constraints on the process of searching, to reduce the size of the search space so that useful results may be achieved with the available computational resources.

2.2 The multiple alignment concept

An example of multiple alignment in the SP system is shown in Figure 1. Here, row 0 contains a New pattern representing a sentence: 't w o k i t t e n s p l a y', while each of rows 1 to 8 contains an Old pattern representing a grammatical rule or a word with grammatical markers. This multiple alignment, which achieves the effect of parsing the sentence in terms of grammatical structures, is the best of several built by the SP62 model when it is supplied with the New pattern and a set of Old patterns that includes those shown in the figure and several others as well. In this example, and others in this article, 'best' means that the multiple alignment in the figure is the one that enables the New pattern to be encoded most economically in terms

[4]With realistically large volumes of data.

of the Old patterns. Details of how the encoding is done may be found in Wolff (2006, Section 3.5).

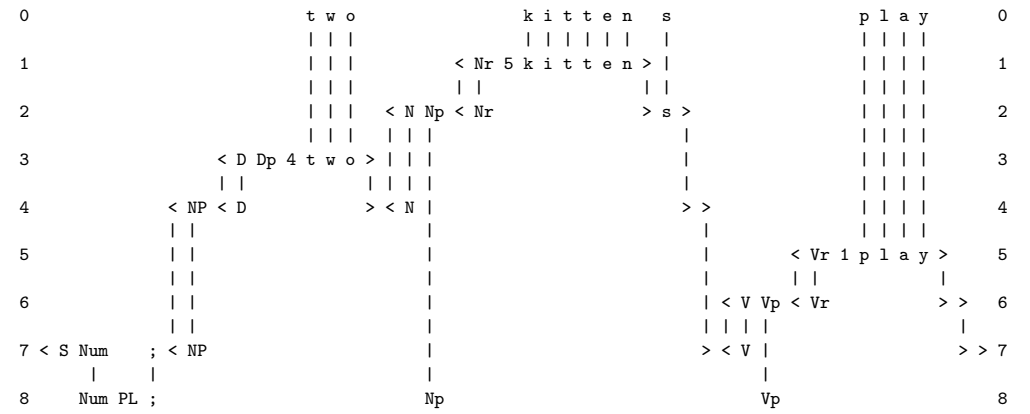

Figure 1: The best multiple alignment created by the SP62 model with a store of Old patterns like those in rows 1 to 8 (representing grammatical structures, including words) and a New pattern (representing a sentence to be parsed) shown in row 0. Reproduced from Figure 1 in Wolff (2007), with permission.

A point of interest about this multiple alignment is the way that, in row 8, the symbols 'Np' and 'Vp' mark the grammatical dependency between the plural subject of the sentence ('k i t t e n s') and the plural main verb ('p l a y'). This kind of dependency is often described as 'discontinous' because there may be arbitrarily large amounts of intervening structure between one element of the dependency and another. This method of marking discontinous dependencies is, arguably, simpler and more elegant than how they are marked in other grammatical systems.

2.2.1 Versatility of the multiple alignment concept

Much of the descriptive and explanatory power of the SP theory is due to the versatility of the multiple alignment concept in:

- *The representation of knowledge.* Despite the simplicity of SP patterns, the way they are processed within the multiple alignment framework gives them the versatility to represent several kinds of knowledge, including grammars for natural languages, ontologies, class hierarchies, part-whole hierarchies, decision networks and trees, relational tuples, if-then rules, associations of medical signs and symptoms, causal relations, and concepts in mathematics and logic such as 'function', 'variable', 'value', and 'set'.

- *The processing of knowledge.* The SP system has demonstrable capabilities in several areas, including natural language processing, pattern recognition, several kinds of reasoning, the storage and retrieval of information, planning, problem solving, unsupervised learning, and information compression.

2.3 Origins of the SP theory

It is pertinent to mention that part of the inspiration for the SP theory is research by Fred Attneave (eg, Attneave, 1954), Horace Barlow (eg, Barlow, 1969), and others, showing that aspects of visual perception (and, more generally, the workings of brains and nervous systems) may be understood in terms of information compression.

Other sources of inspiration for the SP theory include research on 'minimum length encoding' (eg, Solomonoff, 1964), and evidence for the importance of information compression in the unsupervised learning of language (eg, Wolff, 1988),[5] and in mathematics and logic (Wolff, 2006, Chapters 2 and 10).

2.4 Compression, efficiency, and prediction

At an abstract level, information compression brings three main benefits:

- For any given body of information, I, it reduces the amount of storage space required.
- Reducing the size of I can mean increases in efficiency. It would, for example, mean less searching if we are trying to find something within I.
- Perhaps most importantly, information compression provides the key to inductive prediction. In the SP system, it is the basis for all kinds of inference, and for calculations of probabilities.

In animals, we would expect these things to have been favoured by natural selection because of the competitive advantage they can bring. And they are likely to be useful in artificial systems.

In the SP framework, information compression is achieved via the discovery of recurrent patterns (like those shown in rows 1 to 8 in Figure 1 and columns 1 to 6 in Figure 7), and also via the economical encoding of New information in terms of Old patterns, as explained in Wolff (2006, Section 3.5).

[5]Details of other relevant publications may be found via bit.ly/12DOkTV.

3 Low-level perceptual features

It is now widely accepted that, at 'low' levels in vertebrate and invertebrate visual systems, there are processes that recognise perceptual features such as edges and corners. Some relevant evidence is outlined in subsections below.

In this section, the main focus is on features that may be regarded as 'explicit' because they derive directly from visual input. But it is well known that we may 'see' things that have little or no counterpart in the visual input, such as the 'subjective contours' in Marr (2010, Figure 2-6) or the edge of one leaf where it overlaps another in Marr (2010, Figure 4-1 (a)). These kinds of 'implicit' features will be considered in Section 7.1.

In two respects, explicit perceptual features sit comfortably with the SP theory:

- They may be seen to provide a means of encoding perceptual information in an economical manner. For example, Attneave (1954) writes that "Common objects may be represented with great economy, and fairly striking fidelity, by copying the points at which their contours change direction maximally, and then connecting these points appropriately with a straight edge." (p. 185). He illustrates this with the now-famous picture of a sleeping cat, reproduced in Figure 2.

- At lowish levels, perceptual features may function as if they were the atomic symbols that provide the foundation for all higher-level structures, even though they themselves have been constructed from lower-level components.

As just indicated, vision begins with images as they are first projected, not perceptual features. The latter must be somehow discovered or detected within the images. The following subsections consider how the SP theory may be applied in this area, starting with a consideration of options for the encoding of light intensities.

3.1 The encoding of light intensities

In the design of artificial systems for vision, it seems natural and obvious that light intensities in images should be expressed as numbers. But, in itself, the SP system recognises only atomic symbols that can be matched in an all-or-nothing manner with other atomic symbols. It is true that, in principle, it may be supplied with patterns that express Peano's axioms or similar information, and it may then interpret numbers correctly (see Wolff, 2006, Chapter 10). But this has not yet been explored in any depth and,

Figure 2: Drawing made by abstracting 38 points of maximum curvature from the contours of a sleeping cat, and connecting these points appropriately with a straight edge. Reproduced from Figure 3 in Attneave (1954), with permission.

in any case, numbers are probably a distraction in understanding how SP principles may be applied to vision.

To simplify the discussion here, we shall assume that we are processing monochrome images with just two categories of pixel: black and white. With that kind of representation, the lightness in any given small area may be encoded via the *densities* of black and white pixels in that area, without using explicit numbers.[6] It is true that such pixels may be represented with the symbols '1' and '0' but these are simply atomic symbols (as required by the SP system), without numerical meanings.

3.2 Edge detection with neurons

It is relevant to this discussion to consider briefly how edges may be detected with neurons. Figure 3 shows two sets of recordings from a single visual receptor ('ommatidium') of the horseshoe crab, *Limulus*. In both sets of recordings, the eye of the crab was illuminated in a rectangular area bordered by a dark rectangle of the same size (producing a step function as shown at the top right of the figure). In both cases, successive recordings were taken with the pair of rectangles in successive positions across the eye along a line which is at right angles to the boundary between light and dark areas. This achieves the same effect as—but is easier to implement than—keeping the two rectangles in one position and taking recordings from a range of receptors across the light and dark areas.

[6]This is somewhat like the encoding of dark and light in newspaper photographs, at

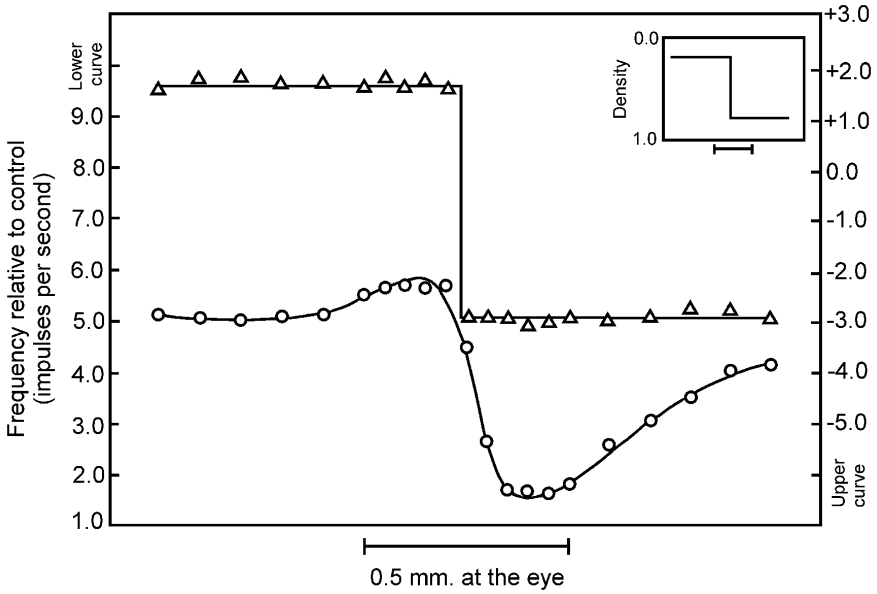

Figure 3: Two sets of recordings from a single ommatidium of *Limulus* (Ratliff and Hartline, 1959, p. 1248). Reproduced from Figure 4, *The Journal of General Physiology, 42*, p. 1248, by copyright permission of The Rockefeller University Press.

In the top set of recordings (triangles) all the ommatidia except the one from which recordings were being taken were masked from receiving any light. In this case, the target receptor responds with frequent impulses when the light is bright and at a sharply lower rate in the dark. In the bottom set of recordings (circles) the mask was removed so that all the ommatidia were exposed to the pattern of light and dark rectangles. In this case, positive and negative responses are exaggerated near the border between light and dark areas but the target receptor fires at or near a background rate in areas which are evenly illuminated (either light or dark). This kind of effect—which is seen elsewhere in the animal kingdom—appears to be due to lateral inhibition between neurons in the visual system (von Békésy, 1967, pp 172-174).

It has been recognised for some time that the dampening of the response in regions of uniform illumination (light or dark) may be seen to achieve the effect of compressing visual information by extracting redundancy from it (Barlow, 1959). It is somewhat like the 'run-length coding' technique for compression of information: a symbol or group of symbols that repeats in

---

least as they used to be.

a contiguous sequence may be reduced to a single instance, perhaps marked for repetition.[7] A boundary between one uniform area and another may be represented economically by two such compressed representations, side-by-side. In the neural case, the upswing near the light/dark boundary may be seen as an economical representation of the idea that the whole of the preceding area is light, the downswing on the other side may be seen as a succinct marking of the fact that the following area is dark, while the two together may be seen to serve as a compressed representation of the boundary.

Although it is less directly relevant to the present discussion, it is pertinent to mention that there are 'complex' cells in mammalian visual systems that respond selectively to edges, and also to 'lines' and 'slits' (see, for example, Frisby and Stone, 2010, pp 215-219).

3.3 Edge detection with the SP system

In the SP framework, the effect of run-length coding may be achieved via recursion, as illustrated in Figure 4.

```
0     a b c     a b c     a b c     a b c                   0
      | | |     | | |     | | |     | | |
1 X 1 a b c X 1 | | |     | | |     | | |             #X #X 1
            | | | | |     | | |     | | |             |
2           X 1 a b c X 1 | | |     | | |           #X #X    2
                      | | | | |     | | |           |
3                     X 1 a b c X 1 | | |        #X #X       3
                                  | | | | |        |
4                                 X 1 a b c X 1 #X #X          4
```

Figure 4: The best multiple alignment produced by the SP62 model with the New pattern '`a b c a b c a b c a b c`' and multiple appearances of the Old pattern, '`X 1 a b c X 1 #X #X`'.

Here, each instance of '`a b c`' in the New pattern in row 0 is matched to an appearance[8] of the self-referential Old pattern '`X 1 a b c X 1 #X #X`'. It is self-referential because '`X 1 #X`' in the body of the pattern may be matched and unified with '`X 1 ... #X`' at the start and end of the pattern.

The encoding of the New pattern which we may derive from this multiple

[7] See, for example, 'Run-length encoding', Wikipedia, bit.ly/eyxlY, retrieved 2013-02-04.

[8] In the SP framework, any Old pattern may appear more than once in a multiple alignment. Here, an *appearance* of a pattern is *not* the same as an *instance* of a pattern, as explained in Wolff (2006, Section 3.4.6).

alignment is the relatively short sequence '`X 1 #X`'.[9] As before, two such encodings, side-by-side, would be an economical representation of the boundary between one uniform region and another.

Of course, this does not look much like lateral inhibition with neurons, as outlined in Section 3.2. But at an abstract level, the two things may be seen to produce the same result: the extraction of redundancy from uniform regions, leaving information about the boundaries between such regions as an economical representation of the raw data, like David Marr's (2010) 'primal sketch'.

With other developments—such as the generalisation of the SP concepts to two dimensions—this kind of technique may be applied in computer vision. Meanwhile, existing techniques, such as those described in Szeliski (2011, Chapter 4), may serve instead.

3.4 Orientations, lengths, and corners

So far, we have said nothing about the orientations of edges or their lengths. In principle, those things may be encoded mathematically, and very economically, in the manner of computer graphics. But that does not seem very likely in a biological system and it is not necessarily the best option for any artificial system that aspires to human-like capabilities in vision.

As mentioned above, the visual cortex in mammals is populated by large numbers of 'complex' neurons, each one of which responds to an 'edge', 'slit', or 'line', at a particular orientation. There is a good coverage of different angles within each small area (see, for example, Frisby and Stone, 2010, Chapter 9). These observations suggests that, in natural vision, the orientation of any edge may be encoded quite simply and directly in terms of the corresponding type of neuron, and likewise in an artificial system.

A sequence of such codes would describe both the orientation and length of a line but it would contain the same kind of redundancy as is discussed in Section 3.3. So we may guess that, in natural vision, some kind of run-length coding may operate, reducing the redundancy within the body of the line and preserving information where the repetition stops—at the points where the line begins and where it ends.

Some relevant evidence comes from studies showing the existence of 'end stopped' hypercomplex cells that respond selectively to a bar of a defined length, or a corner (see, for example, Frisby and Stone, 2010, pp 216-217). In keeping with Attneave's (1954) remarks quoted earlier, we may guess that,

[9] For a description of the method of deriving an encoding from a multiple alignment, see Wolff (2006, Section 3.5).

in mammalian vision, the orientation and length of an edge, slit or line, is to a large extent encoded via neurons that record the beginning and end of the line and any associated corners. Orientation-sensitive neurons would provide the input for this 'higher' level of encoding.

In artificial systems, this kind of coding may in principle be done within the multiple alignment framework, as outlined in Section 3.3. As before, existing techniques may provide stop-gap solutions.

3.5 Noisy data and low-level features

Readers may, with some justice, object that real visual data is rarely as clean as the example in Figure 4 may suggest. Most areas are some shade of grey, not purely black or purely white, and there are likely to be blots and smudges of various kinds.

What appears to be a promising answer to this kind of problem is that the SP system is designed to search for optimal solutions and is not unduly disturbed by errors of omission, commission and substitution. There is more on this topic in Section 4.1 (see also Section 5.7).

4 Object recognition and scene analysis

In some respects, object recognition is like parsing in natural language processing (see, for example, Farabet et al., 2012; Han, 2005). Since the SP system works well in parsing, as outlined in Section 2.2, it may also prove useful in computer vision. Naturally, it would be necessary for the SP machine to have been generalised to work with patterns in two dimensions. And in this discussion we shall assume that low-level perceptual features have been identified, and that they may be treated as atomic symbols, in accordance with the SP theory.

Figure 5 shows schematically how someone's face, with their ears, may be parsed within the multiple alignment framework. Row 0 in the figure contains a New pattern representing incoming information. Each part has been aligned with an Old pattern representing stored knowledge of the structure of an ear, an eye, etc. And these are aligned with a pattern in row 2 representing the higher-level structure of someone's head.

Although this is schematic, I believe the approach has potential, as described in the following subsections.

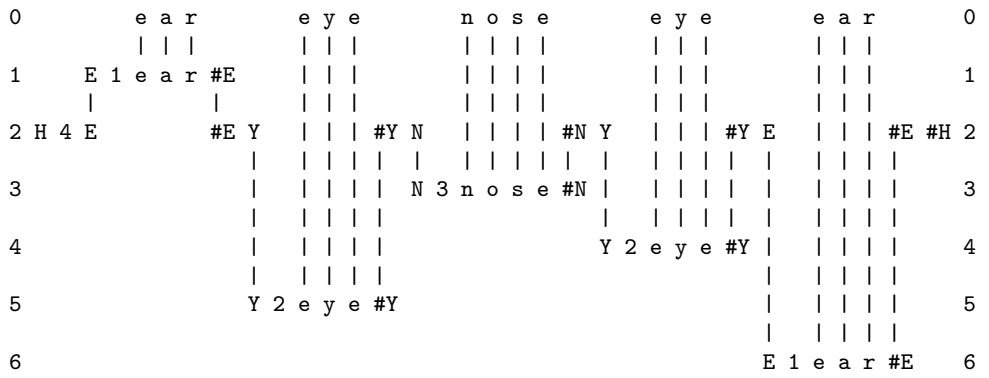

Figure 5: A multiple alignment showing schematically how a person's face, with their ears, may be recognised.

4.1 Noisy data and recognition

Contrary to the impression one might gain from Figure 5, the SP system is quite robust in the face of errors. This is illustrated in Figure 6 where the New pattern in row 0 is the same sentence as in Figure 1 but with the omission of the 'w' in 't w o', the substitution of 'm' for ''n' in 'k i t t e n s', and the addition of 'x' within the word 'p l a y'. Despite these errors, the best multiple alignment created by the SP62 model is, as shown, the one that we judge intuitively to be 'correct'.

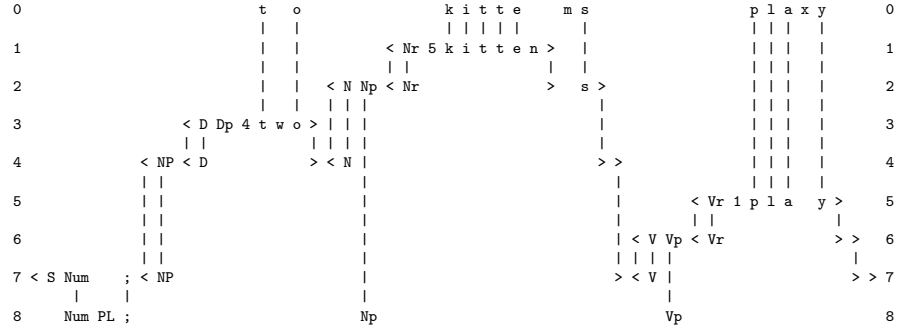

Figure 6: The best multiple alignment created by the SP62 model with a New pattern (row 0) like the one shown in Figure 1 but with errors of omission, commission and substitution, and with same set of Old patterns as before. Reproduced from Figure 2 in Wolff (2007), with permission.

This kind of ability to cope gracefully with noisy data is really essential in any system which aspires to explain or emulate our ability to recognise things despite fog, snow, falling leaves, or other things that may obstruct our view.

In general terms, the reason that the SP models can cope with noisy data is that they search for optimal solutions, without relying on the presence or absence of any particular feature or combination of features.

4.2 Part-whole hierarchies, class hierarchies, and their integration

A strength of the multiple alignment concept is that it provides a simple but effective vehicle for the representation and processing of part-whole hierarchies, class hierarchies, and their integration.

Recognition of an entity in terms of its parts is illustrated rather simply in Figure 5 and more realistically in Figure 1. In the latter case, the sentence is divided into a noun phrase and a main verb, the noun phrase is divided into a determiner and a noun, and the noun contains the root or stem, 'k i t t e n', with the plural suffix, 's'.

Continuing with the feline theme but not illustrated here is the way that, in the multiple alignment framework, a cat may be recognised at several levels of abstraction: as an animal, as a mammal, as a cat, and as a specific individual, say 'Tibs' (Wolff, 2006, Figure 6.7). The framework also provides for the representation of heterarchies or cross classification: a given entity, such as 'Jane' (or a class), may belong in two or more higher-level classes that are not themselves hierarchically related, such as 'woman' and 'doctor'.[10]

The way that part-whole relations and class-inclusion relations may be combined in one multiple alignment is illustrated in Figure 7.[11] Here, some features of an unknown plant are expressed as a set of New patterns, shown in column 0: the plant has chlorophyll, the stem is hairy, it has yellow petals, and so on.

From this multiple alignment, we can see that the unknown plant is most likely to be the Meadow Buttercup, *Ranunculus acris*, as shown in column 1. As such, it belongs in the genus *Ranunculus* (column 6), the family *Ranunculaceae* (column 5), the order *Ranunculales* (column 4), the class *Angiospermae* (column 3), and the phylum *Plants* (column 2).

Each of these higher-level classifications contributes information about

[10] Although the term 'heterarchy' is not widely used, in can be useful as a means of referring to hierarchies in which, as in the example in the text, a given node may appear in two or more higher-level nodes that are not themselves hierarchically related. In the SP framework, there may be heterarchies in both class-inclusion structures and part-whole structures. But to avoid the clumsy expression 'hierarchy or heterarchy', the term 'hierarchy' is used, in most parts of this article, as a shorthand for both concepts.

[11] Compared with multiple alignments shown above, this has been rotated through 90° so that it fits more easily on the page.

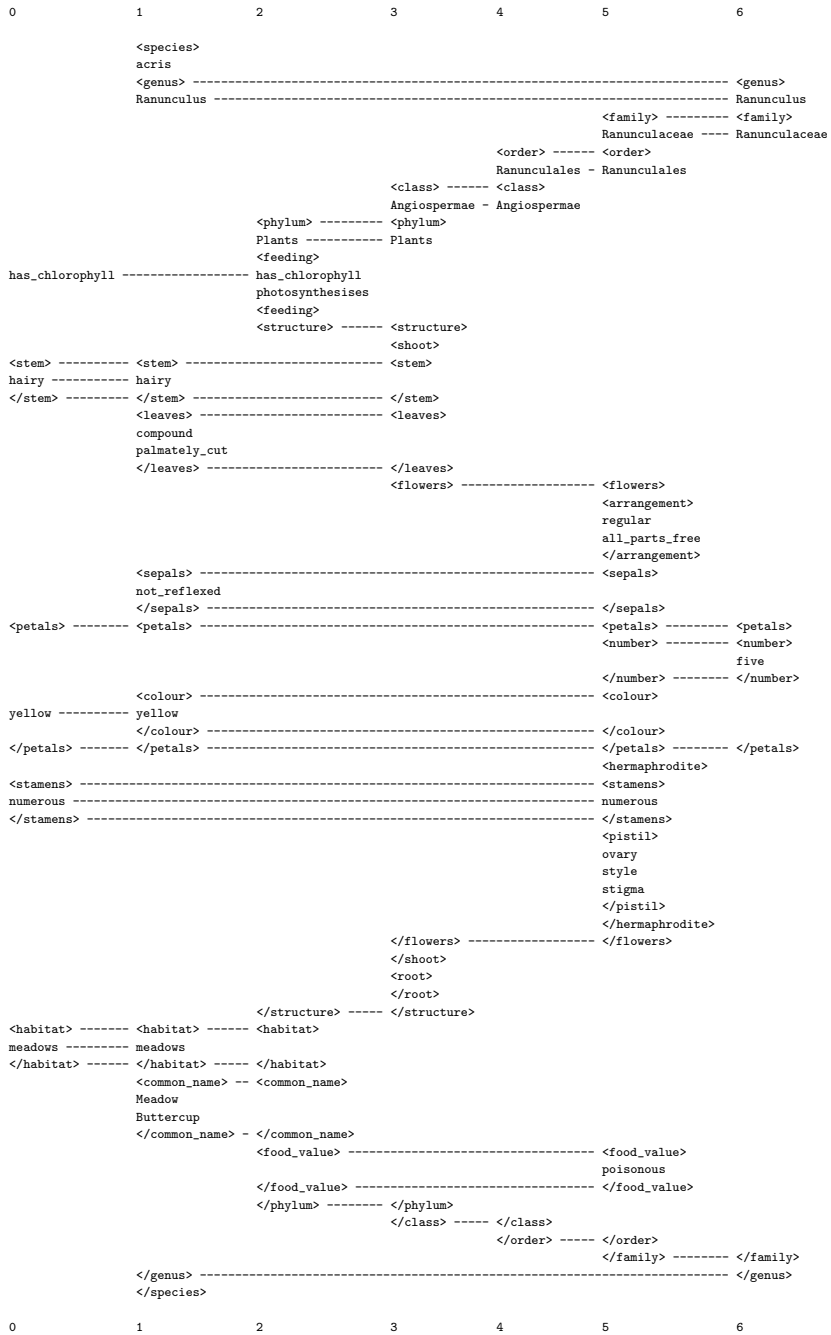

Figure 7: The best multiple alignment created by the SP62 model, with a set of New patterns (in column 0) that describe some features of an unknown plant, and a set of Old patterns, including those shown in columns 1 to 6, that describe the attributes of different categories of plant.

attributes of the plant and its division into parts and sub-parts. For example, as a member of the the class *Angiospermae* (column 3), the plant has a shoot and roots, with the shoot divided into stem, leaves, and flowers; as a member of the family *Ranunculaceae* (column 5), the plant has flowers that are 'regular', with all parts 'free'; as a member of the phylum *Plants* (column 2), the buttercup has chlorophyll and creates its own food by photosynthesis; and so on.

Of course, this example does not describe the visual appearance of an object. But it should be apparent that this system, when it has been generalised to work with patterns in two dimensions, has potential as a means of representing and processing both the parts and sub-parts of an object's image, and how that information relates to any hierarchy of classes to which that object belongs. And each of those two types of hierarchy is a very effective means of expressing visual information in a compressed form.

4.3 Scene analysis

Scene analysis may also be viewed as a kind of parsing (see, for example, Shi, 1983). For the analysis of a seascape, for example, there may be a high-level structure recording the kinds of things that one sees in a typical seascape (sea, beach, rocks, boats, and so on), with a more detailed description for each one of those things.

There seem to be two main complications in scene analysis:

- Any one thing may be partially obscured by another. In our seascape, a boat may be partially obscured by, for example, waves, sea birds, or members of the crew.
- The locations of things may be quite variable. A boat may be in the sea or on the beach; people can appear almost anywhere; and so on.

Of course, people cope easily with both those things, but there may be a problem with 'naive' kinds of parsing system.

The SP framework may accommodate these aspects of scene analysis in three main ways:

- As we saw in Section 4.1, parsing can be done successfully despite errors or omission, commission, or substitution. Thus there is reason to believe that, when the SP models have been generalised to work with patterns in two dimensions, an object may be recognised even if it is partially obscured.

- The variability of scenes is broadly similar to the variability of sentences in natural language. Artificial parsing systems, including the SP system, can cope with that variability by providing information about a wide variety of types of sentences and phrases, including recursive forms such as *This is the man all tattered and torn that kissed the maiden all forlorn that milked the cow with the crumpled horn ....* The same principles may be applied to vision.

- Where existing knowledge can't cope, the system may learn—as discussed in Section 5.2, next.

5 Unsupervised learning and the discovery of objects and classes

It is clear that learning is an integral part of vision since vision is an important means of gaining new information about the world. And it is clear that, in general, we learn via vision in a manner that is 'unsupervised' in the sense that it does not require the intervention of a 'teacher', or the provision of 'negative' samples, or the grading of samples from simple to complex (*cf.* Gold (1967)). We take in information through our eyes (and other senses) and try to make sense of it as best we can.

In this section, we consider unsupervised learning as it has been developed in the SP framework, and how it may be applied in vision. But as background for what follows we first look at the 'DONSVIC' principle in unsupervised learning.

5.1 The discovery of natural structures via information compression (DONSVIC)

In our dealings with the world, certain kinds of structures appear to be more prominent and useful than others: in natural languages, there are words, phrase and sentences; we understand the visual and tactile worlds to be composed of discrete 'objects'; and conceptually, we recognise classes of things like 'person', 'house', 'tree', and so on.

It appears that these 'natural' kinds of structure are significant in our thinking because they provide a means of compressing sensory information, and that compression of information provides the key to their learning or discovery. At first sight, this looks like nonsense because popular programs for compression of information, such as those based on the LZW algorithm, or programs for JPEG compression of images, seem not to recognise anything

resembling words or objects. But those programs are designed to work fast on low-powered computers. With other programs that are slower but more thorough, natural structures can be revealed:

- Figure 8 shows part of a parsing of an unsegmented sample of natural language text created by the MK10 program (Wolff, 1977) using only the information in the sample itself and without any prior dictionary or other knowledge about the structure of language. Although all spaces and punctuation had been removed from the sample, the program does reasonably well in revealing the word structure of the text. Statistical tests confirm that it performs much better than chance.

- The same program does quite well—significantly better than chance—in revealing phrase structures in natural language texts that have been prepared, as before, without spaces or punctuation—but with each word replaced by a symbol for its grammatical category (Wolff, 1980). Although that replacement was done by a person trained in linguistic analysis, the discovery of phrase structure in the sample is done by the program, without assistance.

- The SNPR program for grammar discovery (Wolff, 1982) can, without supervision, derive a plausible grammar from an unsegmented sample of artificial language, including the discovery of words, of grammatical categories of words, and the structure of sentences.

A key feature of both the MK10 program and the SNPR program is compression of information by the matching and unification of patterns. But much the same can be said of ordinary 'utility' programs for data compression. What is distinctive about the MK10 and SNPR programs is that they are designed to search through what is normally a wide variety of alternative ways in which patterns may be matched and unified, and to select those patterns or sets of patterns that yield relatively high levels of compression.

It seems likely that the principles that have been outlined in this subsection may be applied not only to the discovery of words, phrases and grammars in language-like data but also to such things as the discovery of objects in images, and classes of entity in all kinds of data. These principles may be characterised as 'the discovery of natural structures via information compression', or 'DONSVIC' for short.

5.2 Unsupervised learning in the SP system

Although the SP theory has grown out of my earlier work on the unsupervised learning of language, the MK10 and SNPR models are not well suited to the

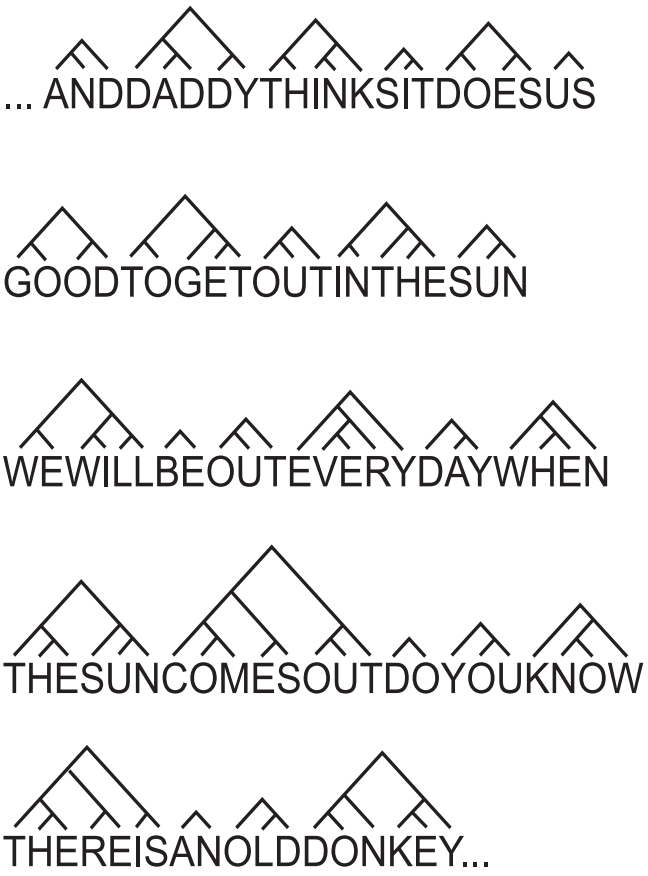

Figure 8: Part of a parsing created by program MK10 (Wolff, 1977) from a 10,000 letter sample of English (book 8A of the Ladybird Reading Series) with all spaces and punctuation removed. The program derived this parsing from the sample alone, without any prior dictionary or other knowledge of the structure of English. Reproduced from Figure 7.3 in Wolff (1988), with permission.

goal of simplifying and integrating concepts across several different aspects of intelligence. It has been necessary to develop a radically new conceptual framework, with the SP concept of multiple alignment at centre-stage. But information compression and the DONSVIC principles are as important in the new conceptual framework as they were before.

As mentioned in Section 2.1, the SP70 model works by creating multiple alignments, deriving Old patterns from the multiple alignments, evaluating sets of newly-created Old patterns in terms of their effectiveness for the economical encoding of the New information, and weeding out low-scoring sets.

The first two of those processes is illustrated schematically in Figure 9. As mentioned earlier, the SP system is conceived as an abstract system that, like a brain, may receive 'New' information via its senses and store some or all of it as 'Old' information. We may think of the 'brain' as that of a baby listening to what people are saying. Let's imagine that he or she hears someone say "t h a t b o y r u n s".[12] If the baby has never heard anything similar, then, if it is stored at all, that New information may be stored as a relatively straightforward copy, something like the Old pattern shown in row 1 of the multiple alignment in part (a) of the figure.

```
0     t h a t g i r l r u n s    0
      | | | |         | | | |
1 A 1 t h a t b o y   r u n s #A 1

(a)

B 2 t h a t #B
C 3 b o y #C
C 4 g i r l #C
D 5 r u n s #D
E 6 B #B C #C D #D #E

(b)
```

Figure 9: (a) A simple multiple alignment from which, in the SP70 model, Old patterns may be derived. (b) Old patterns derived from the multiple alignment shown in (a). Adapted from Figures 9.2 and 9.3 in Wolff (2006), with permission.

Now let us imagine that the information has been stored and that, at some later stage, the baby hears someone say "t h a t g i r l r u n s". Then, from that New information and the previously-stored Old pattern, a multiple alignment may be created like the one shown in part (a) of Figure 9. And, by picking out coherent sequences that are either fully matched or not matched at all, four putative words may be extracted: 't h a t', 'b o y', 'g i r l', and 'r u n s', as shown in the first four patterns in part (b) of the figure. In addition, a fifth pattern may be created, as shown in the figure, that records the sequence 't h a t ... r u n s', with the category 'C #C' in the middle representing a choice between 'b o y' and 'g i r l'. This is the beginnings of a grammar to describe that kind of phrase.

[12]In this and other examples in this subsection, we shall assume that letters are analogues of low-level perceptual features in speech, such as formant ratios or formant transitions.

This example shows how Old patterns may be derived from a multiple alignment but it gives a highly misleading impression of how the SP70 model actually works. In practice, the program forms many multiple alignments that are much less tidy than the one shown and it creates many Old patterns that are clearly 'wrong'. However, the program contains procedures for evaluating candidate sets of patterns and weeding out those that score badly in terms of their effectiveness for encoding the New information economically. Out of all the muddle, it can normally abstract one or two 'best' grammars and these are normally ones that appear intuitively to be 'correct', or nearly so.

As was mentioned in Section 2.1, the SP70 model has two main weaknesses at it stands now: it does not learn intermediate levels in a grammar or discontinuous dependencies of the kind mentioned in Section 2.2. But I believe some reorganisation of the model would solve both problems and greatly enhance the model's capabilities.

5.3 The discovery of objects via stereo matching

As with the structures of natural language, it is clear that we have to learn the structures that are significant in vision, including objects.[13] Some insights into how this may be done may be gained from a consideration of random-dot stereograms like the one shown in Figure 10.

Here, each of the two images is a random array of black and white pixels, with no discernable structure. But there is a relationship between them, as shown in Figure 11: both images are the same except that a square area near the middle of the left image is further to the left in the right image.

When these images are viewed in a stereoscope, the central square appears as a discrete object suspended above the background.[14] The focus of interest here will be on how we come to see that discrete object, while possible implications for our understanding of depth perception are discussed in Section 6.

A little analysis shows that seeing the central square means finding an alignment between pixels in the left image and pixels in the right image, that there are many alternative such alignments, and that some are better than others. One solution is the algorithm developed by Marr and Poggio (1979).

[13]The Chomskian doctrine that children are born with a knowledge of 'universal grammar' fails to account for the specifics of syntactic forms in different languages, and it depends on the still-unproven idea that there is something of substance that is shared by all the world's languages.

[14]Some people are able to see the square by viewing the images directly, with some defocussing to help merge them into one.

Figure 10: A random-dot stereogram from Julesz (1971, p. 21), reproduced with permission of Lucent Technologies Inc./Bell Labs.

1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	Y	*A*	*A*	*B*	*B*	0	0
1	1	1	X	*B*	*A*	*B*	*A*	0	1
0	0	1	X	*A*	*A*	*B*	*A*	1	0
1	1	1	Y	*B*	*B*	*A*	*B*	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

1	0	1	0	1	0	0	1	0	1
1	0	0	1	0	1	0	1	0	0
0	0	1	1	0	1	1	0	1	0
0	1	0	*A*	*A*	*B*	*B*	X	0	0
1	1	1	*B*	*A*	*B*	*A*	Y	0	1
0	0	1	*A*	*A*	*B*	*A*	Y	1	0
1	1	1	*B*	*B*	*A*	*B*	X	0	1
1	0	0	1	1	0	1	1	0	1
1	1	0	0	1	1	0	1	1	1
0	1	0	0	0	1	1	1	1	0

Figure 11: Diagram to show the relationship between the left and right images in Figure 10. Reproduced from Julesz (1971, p. 21), with permission of Lucent Technologies Inc./Bell Labs.

Another solution, potentially, is the kind of processing that builds multiple alignments in the SP models, but generalised for two dimensions. As noted in Section 2.1.1, the complexity of the matching problem can, in general, be reduced by applying constraints to the process of searching and thus reducing the size of the search space.

Figure 12 shows how the SP62 model can solve a one-dimensional analogue of the stereo matching problem. Here, the Old pattern (row 1) may be seen as an analogue of the left image and the New pattern (row 0) may be

seen to stand in for the right image. Both patterns have been prepared from a random sequence of digits,[15] with a displacement of the middle section, much as in Figure 11. This multiple alignment is the best of several different multiple alignments created by the SP62 model with those two patterns.

```
0     4 7 4 6 4 1 3 7 5     8 5 2 4 0 2 9 1 9 3 8 0 1 4 1 1 2 9 7 1 2    0
      | | | | | | | | |     | | | | | | | | | | |     | | | | | | | | |
1 J a 4 7 4 6 4 1 3 7 5 9 4 8 5 2 4 0 2 9 1 9 3     1 4 1 1 2 9 7 1 2 #J 1
```

Figure 12: The best multiple alignment created by SP62 with an Old pattern (row 1) and a New pattern (row 0) as one-dimensional analogues of the left and right images in a random-dot stereogram.

In the figure, one can see how the central sequence of 10 integers (analogous to the central square in Figure 11) has been isolated from the 'background' sequences to the left and right, and this despite repetitions of integers in both patterns and the formation of plenty of 'wrong' alignments on the route to the 'correct' result. It seems likely that the processes can be generalised to work with patterns in two dimensions.

5.4 Structure from motion

The kinds of processing just described may also be applied to objects in motion.

Consider, for example, a flatfish with a sandy, speckled colouration, lying on a sandy and speckled area on the bed of the sea. Such a creature would be very well camouflaged but with one proviso: it must stay still. As soon as it moves, it will become very much easier to see. Why? Apart from the motion itself, an important reason seems to be that movement creates two images (or more), rather like the two images in a random-dot stereogram. And by a process of matching, much as described above, a predator or other observer will be able to see the fish standing out as a distinct entity with distinct boundaries—like the square that can be seen when the two images in Figure 10 are viewed in a stereoscope.

More generally, we see any object in motion—such as a car travelling along a road—as a single entity, not a multitude of images like the frames in a video or film. In all such cases, we merge the many instances into one. The process of merging those many instances, which is likely to yield high

[15]The random sequence of digits, with values between 0 and 9, inclusive, was generated by the Random Integer Generator from Random.org (www.random.org). The results are, they say, better than with pseudo-random number algorithms because "atmospheric noise" is the source of randomness.

levels of compression, requires a process of matching and unification, much as before. And those processes serve to define the boundaries of the entity and to distinguish it from the background.

5.5 Deriving concepts from fragments

If we only ever see parts of an object—perhaps a rare creature in its natural habitat that we have only seen in fleeting glimpses—we can nevertheless develop a coherent concept of the whole object via alignments amongst the fragmentary views: '`A B`' may be aligned with '`B C`' and unified to create '`A B C`'; '`C D`' may be aligned with '`D E`' to create '`C D E`'; '`A B C`' may be aligned with '`C D E`' ..., and so on. This is like the 'sequence assembly' technique in bioinformatics,[16] or the stitching together of overlapping photos to create a panorama. And the matching may be achieved via multiple alignment, as developed in the SP theory.

5.6 The discovery of classes of entity

Similar things may be said about the learning of everyday concepts like 'person' or 'house', or the more formal botanical categories shown in Figure 7. If, for example, we see one thing with the characteristics '`A B C f L M N p X Y Z`' and another with the characteristics '`A B C g L M N q X Y Z`', we may create a unified pattern like this: '`A B C 1 #1 L M N 2 #2 X Y Z`', with the patterns '`1 f #1`', '`1 g #1`', '`2 p #2`', and '`2 q #2`', to fill in the slots. The unified pattern may be seen to represent the class of things with the characteristics '`A B C ... L M N ... X Y Z`'.

This example is, of course, rather similar to the example shown in Section 5.2. That similarity is not accidental. It derives from the principle, which is a key part of the SP theory, that, with compression of information via the multiple alignment framework, all kinds of knowledge may be represented economically with SP patterns. And it is consistent with the long-established idea that there may be a syntax for images, not just natural languages (see, for example, Fu, 1977), and with the previously-mentioned idea that object recognition and scene analysis may each be seen as a form of parsing (Section 4).

There is potential with this kind of learning to create structures that are quite subtle and expressive. Despite its limitations, the SP70 model can already discover grammatical structures with alternatives everywhere, and without any fixed elements as in '`A B C ... L M N ... X Y Z`'. It is

[16] See en.wikipedia.org/wiki/Sequence_assembly, retrieved 2013-02-21.

envisaged that, with the kind of reorganisation mentioned earlier, the system should be able discover structures that express part-whole hierarchies and class-inclusion hierarchies, both of them with multiple levels, and to abstract discontinuous dependencies in data of the kind mentioned in Section 2.2.

5.7 Noisy data and learning

As was noted in Sections 3.5 and 4.1, visual information is normally 'noisy' in the sense that, compared with any stored information, it is likely to contain errors of omission, commission, or substitution, in any combination. As shown in Figure 6, the SP system has a capacity to cope with these kinds of errors, at least in tasks like parsing, recognition, or scene analysis.

What about learning? How can any system learn 'correct' structures from noisy data in an 'unsupervised' manner and without any help from a 'teacher', or from examples that are marked as 'wrong', or from anything else of that kind? This is not merely an issue in vision. It also arises in connection with language learning, as illustrated in Figure 13.

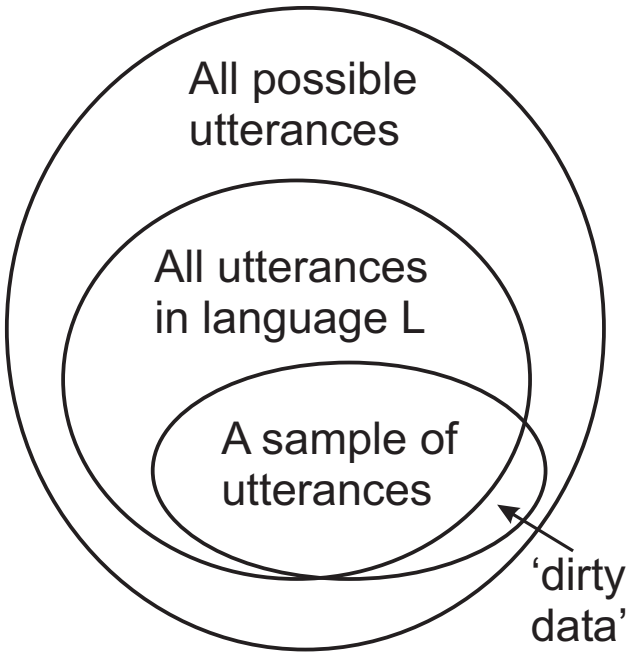

Figure 13: Categories of utterances involved in the learning of a first language, L. In ascending order size, they are: the finite sample of utterances from which a child learns; the (infinite) set of utterances in L; and the (infinite) set of all possible utterances. Adapted from Figure 7.1 in Wolff (1988), with permission.

When we learn our first language or languages, we learn from what we hear—a finite sample of language shown as the smallest envelope in the figure. But there are two apparent problems:

- How we learn despite what is marked in the figure as 'dirty data': sentences that are not complete, false starts, words that are mis-pronounced, and more.

- How we generalise from the finite sample represented by the smallest envelope to a knowledge of the language corresponding to the middle-sized envelope, without overgeneralising into the region between the middle envelope and the outer one.

One possible answer is that mistakes are corrected by parents, teachers, and others. But the weight of evidence is that children can learn their first language without that kind of assistance.[17]

An alternative answer favoured here is that information compression provides the key:

- Any particular error is, by its nature, rare and so in the search for useful patterns (which, other things being equal, are the more frequently-occurring ones), it is discarded along with many other candidate structures.[18]

- As a general rule, the highest levels of compression can be achieved with grammars that represent moderate levels of generalisation, neither too little nor too much.[19]

In practice, the MK10 and SNPR programs have been found to be quite insensitive to errors (of omission, addition, or substitution) in their data. And the SNPR program has been shown to produce plausible generalisations, without over-generalising (Wolff, 1988).

Since the principles are general, it seems likely that visual learning within the SP framework may be achieved in the face of noisy data.

[17]Relevant evident comes from cases where children learn to understand language even though they have little or no ability to speak (Lenneberg, 1962; Brown, 1973)—so that there is little or nothing for anyone to correct.

[18]If an error is not rare it is likely to acquire the status of a dialect or idiolect variation and cease to be regarded as an error.

[19]Notice that this principle applies to lossless compression as well as lossy compression.

6 Space and depth

As mentioned earlier, it is envisaged that, in the SP theory, all kinds of knowledge will be represented with patterns in one or two dimensions. Superficially, this seems to rule out anything with more dimensions, and suggests that there might be a need to introduce patterns with three dimensions and possibly more. However, this has been rejected, at least for the time being, for these main reasons:

- Although the multiple alignment concept may in principle be generalised to patterns in three or more dimensions, it is difficult to see how it could be made to work in practice and it looks implausible as a model for any kind of structure or process in the brain.
- A tentative part of the SP theory is the idea that the cortex of the brains of mammals—which is, topologically, a two-dimensional sheet—may be, in some respects, like a sheet of paper on which 'pattern assemblies' (neural analogues of SP patterns) may be written (Wolff, 2006, Chapter 11)—as shown schematically in Figure 14.
- If we exclude processes of interpretation in terms of harmonics, colours, or the like, raw sensory data may be seen to come in either one dimension (eg sound) or two (eg visual images).
- Three-dimensional structures may be represented with patterns in two dimensions, somewhat in the manner of architects' drawings (Wolff, 2006, Section 13.2.2). With the development of mathematical concepts within the SP framework (Wolff, 2006, Chapter 10), four or more dimensions may be represented in much the same way as is done now with mathematical techniques.

6.1 Three-dimensional objects

This and the following two subsections consider some aspects of the visual perception of space and depth, and whether or how the SP theory may be applied.

If an object is viewed from several different angles, with overlap between one view and the next (as illustrated in Figure 15), the several views may be stitched together to create what is at least a partial and approximate 3D model of the object. This is similar to the piecing together of fragments to create a coherent concept, as outlined in Section 5.5. As before, it may be

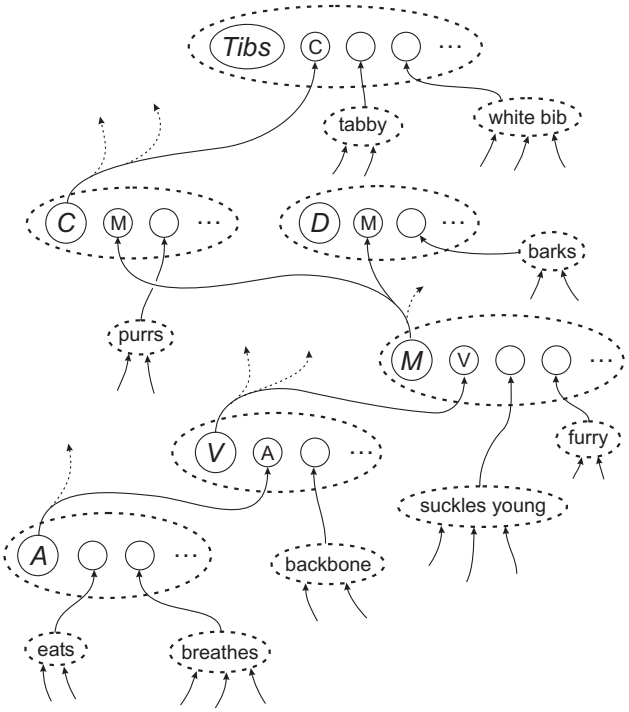

Figure 14: Schematic representation of hypothesised neural analogues of SP patterns and their inter-connections. *Key*: 'C' = cat, 'D' = dog, 'M' = mammal, 'V' = vertebrate, 'A' = animal, '...' = further structure that would be shown in a more comprehensive example. Pattern assemblies are surrounded by broken lines and each neuron is represented by an unbroken circle or ellipse. Lines with arrows show connections between pattern assemblies and the flow of sensory signals. Connections between neurons within each pattern assembly are not marked. The figure is reproduced from Figure 11.6 of Wolff (2006), with permission.

achieved via multiple alignment as that concept has been developed in the SP theory.

The model will be partial if, for example, it excludes views from above or below. And it is likely to be approximate because a given set of views may not be sufficient for an unambiguous definition of the object's geometry: there may be variations in the shape that would be compatible with the given set of views.

Do these deficiencies matter? For many practical purposes, the answer is likely to be "no". If we want a rock to put in a rockery, or a stick to throw for a dog, the exact shape is not important. And if we want more accurate information, we can inspect the object more closely, or supplement vision

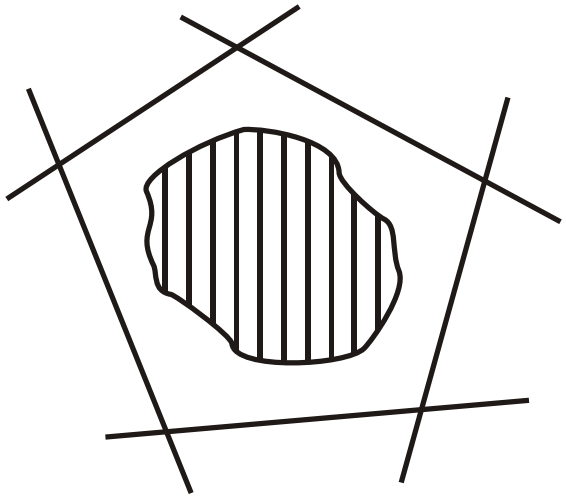

Figure 15: Plan view of a 3D object, with each of the five lines around it representing a view of the object, as seen from the side.

with touch.

Evidence that people do something like what has been described is our ordinary experience that things can be harder to recognised from unfamiliar viewpoints than from familiar ones—the basis of some trick photos. That observation is confirmed in experimental studies showing that people are both slower at recognising things, and less accurate, when the viewpoint is unfamiliar (Tarr, 1995; Bülthoff and Edelman, 1992; Tarr and Pinker, 1989).

Although what has been described is like the stitching together of overlapping photos to create a panorama, the SP theory suggests that, with people, the visual information would be compressed via the encoding, within the SP system, of part-whole relations, class-inclusion relations, and other kinds of regularities.[20] That compression can be of benefit in both natural and artificial systems, as indicated in Section 2.4.

6.2 Building a model of one's environment and finding one's way around

Similar processes may be at work when we move around in our environment and learn about it. Successive views that overlap each other may be stitched

[20]It is true that, in digital systems, photos are normally compressed via JPEG or similar technique. But, as indicated in Section 5.1, there is potential for the SP system to yield higher levels of compression and more natural structures.

together, as before, to create a model of the streets or other places where we have been. This is essentially what has been and is being done with Google's 'Street View'.[21] The main difference between what has been achieved with Street View and what is envisaged for the SP system is that, in the latter case, visual information would be compressed via the mechanisms in the SP system, as noted in Section 6.1.

As with objects (Section 6.1), a model of our environment that is created via overlapping views may not be geometrically precise.[22] But, as before, some ambiguity may not matter very much for many practical purposes. Topological maps, such as the classic map of the London underground, can be quite good enough for finding one's way around. However, if greater geometric accuracy is needed, it may be increased by gathering more information, especially information about areas between roads, paths or other routes.

In connection with finding one's way around, the SP system may be relevant in two ways:

- If a robot has stored representations of one or more places, perhaps compressed via recurrent patterns as indicated in Section 2.4, then, via the building of multiple alignments (as in Section 4), it should be able to recognise when it has reached one of those places, using incoming visual information as New patterns and stored knowledge as Old patterns. If it has stored information about an entire route or network of routes, then, within that environment, it should be able to identify where it is at any time. Similar things may be true of people.

- With an appropriate set of Old patterns, each one of which represents a direct connection between two places, the SP system, via the building of multiple alignments, can work out one or more routes between any two of the relevant places, including routes via two or more of the direct connections (Wolff, 2006, Chapter 8). The example in Figure 16 shows one such flying route between Beijing and New York.

These points about how we may build a model of our environment and find our way around relate to the topic of 'simultaneous localization and mapping' (SLAM) in robotics.[23]

[21] See maps.google.com/help/maps/streetview/.

[22] From my own experience of exploring caves, I know that, while one can build up a good knowledge of how different passages connect with each other, one's understanding of their 3D geometry can be hazy, and can lead to some surprises if one has the opportunity to see a 3D model that is based on a proper survey, with measurements of distances and angles.

[23] See, for example, en.wikipedia.org/wiki/Simultaneous_localization_and_mapping,

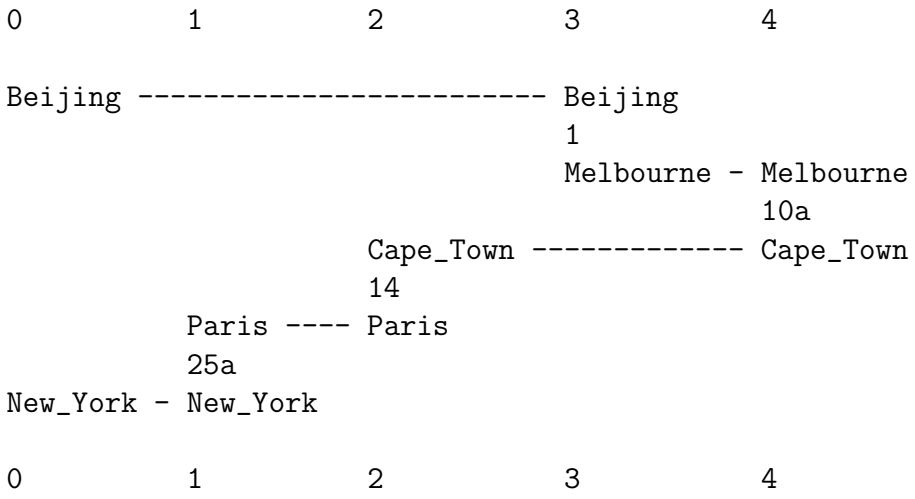

Figure 16: A multiple alignment showing a flying route between Beijing and New York, one of several produced by the SP61 model with a set of Old patterns, one for each leg of this and other possible journeys. Reproduced from Figure 8.5 (e) in Wolff (2006), with permission.

6.3 Depth perception and stereoscopic vision

Without attempting a comprehensive discussion of the complex subject of depth perception, this section offers some thoughts about stereoscopic vision, and the possible relevance of the SP theory.

6.3.1 Triangulation

For any given object that we are looking at, we can in principle work out its distance by a process of triangulation like that which has been widely used in cartography, at least as it used to be. But there appear to be snags:

- For this mechanism to work with reasonable accuracy, it would be necessary for one to have a rather accurate sense of the direction of gaze for each eye and the angle between that direction of gaze and the line between the two eyes. It seems unlikely that we can sense the positions of our eyes with the necessary accuracy.
- There is evidence that, with the Ames' distorted room illusion,[24], the illusion persists when people view the room with two eyes (Glenner-

retrieved 2013-03-07.

[24]For readers who are not familiar with this illusion, a person looks into one end of a room that appears to have a conventional rectangular form but is actually constructed so that one of the two corners opposite the viewer is stretched away and is relatively high,

ster et al., 2003), although, in that case, the effect may be reduced (Gehringer and Engel, 1986). This suggests that any information about distance that may be gained via triangulation[25] is not sufficiently clear or precise to overcome viewers' preconceptions that the room has the conventional rectangular form.

- Triangulation cannot work with a stereoscope or a 3D film because what we are looking at is all at one distance, with nothing to differentiate one part of the picture from another. The spear which makes us jump as we see it coming towards us out of a 3D film is no closer to us than anything else in the film.

We cannot rule out triangulation altogether—it may have a role in some situations—but some other mechanism is needed to explain how we see depth with a stereoscope or a 3D film.

6.3.2 Possible alternatives

With random-dot stereograms, it is clear that our brains are capable of forming an alignment between the left and right images that is good enough to identify the displaced area in the middle as a discrete entity (Section 5.3). By identifying the displaced area and distinguishing it from the surrounding area, we may also gain an accurate knowledge of the size of the displacement.

How can the size of the displacement tell us about depth? There are at least three possible answers (which are not necessarily mutually exclusive):

- For any given displacement, our brains perform a geometrical calculation of what that displacement implies about relative distances, between the observer and the perceived object, and between the perceived object and the background.

- We are born with knowledge that is, in effect, a table of associations between displacements and distances.

- We learn those kinds of associations from experience.

That learning is important is suggested by the powerful influence of our experience (of rectangular rooms) in the Ames' room illusion. Building up a knowledge of associations is part of what the SP system is designed to achieve.

---

while the other corner is nearer to the viewer and is relatively low. Anyone standing in the near corner appears to be large, and they appear to shrink if they walk to the far corner.

[25] Or any other clue such as the focussing of our eyes.

7 Some other aspects of vision

The SP theory has things to say about some other aspects of vision, as discussed in the following subsections.

7.1 Seeing things that are not there

As noted in Section 3, we often 'see' things that are not objectively present in what we are looking at. We may see 'subjective contours' in certain kinds of images, or we may see the edge of a leaf where it overlaps another leaf despite there being little or nothing to mark the boundary.

The multiple alignment in Figure 1 provides an example of how the SP system may accommodate these kinds of things. Here, the New pattern is the sentence 't w o k i t t e n s p l a y' with nothing to mark the boundary between one word and the next. But those boundaries are clearly marked via the parsing of the sentence into its constituent parts.

More generally, we infer things that are not immediately visible: when we see the unbroken shell of a hazel nut, we expect to find an edible kernel inside; when we see a horse partially obscured by a tree, we expect to see the whole animal when it moves into full view; and so on. This kind of inference is an integral part of how the SP system works. In Figure 6, the word 't w o' appears in the New pattern as 't o', but the parsing interpolates the missing 'w'. In Figure 7, the rather sketchy information in column 1 is extended via the information in columns 1 to 6: we can infer that the plant photosynthesises (column 2), that it has five petals (column 6), that it is poisonous (column 5), and so on.

7.2 Recognition despite variations in image size

A prominent feature of natural vision is that we can recognise something despite wide variations in viewing distance and corresponding variations in the size of the retinal image.[26] Although this phenomenon is not consistent with any simple pattern-matching model of vision, it appears that it can be accommodated within the SP theory.

Let us suppose that, as described in Section 3.3, the image to be processed is reduced to a 'primal sketch', showing boundaries between uniform areas but without the redundancy within those areas. For any given scene, the effect of that processing will be to reduce or eliminate variations in the size

[26]This is related to but different from the phenomenon of 'size constancy': that, within limits, our perception of an objects size remains the same, regardless of viewing distance or the size of the retinal image.

of the original image. The primal sketch that is derived from a large version of the scene will be much the same as the primal sketch that is derived from a small version.

Any residual variations in size, or noise in the image, may be overcome by the flexibility of the matching process in the SP system (Section 2.1) and by the system's ability to tolerate noise (Sections 3.5, 4.1, and 5.7).

7.3 Lightness constancy and colour constancy

Another prominent feature of natural vision is 'lightness constancy': the fact that, normally, we perceive the lightness of an object to be fixed, despite wide variations in the incident light and corresponding variations in the amount of light that is reflected from the object (its 'luminence'). We would normally see a lump of coal as black and snow as white, even though the coal in bright sunlight may be reflecting more light per unit area than snow in shadow.

In order to account for this phenomenon, it seems necessary to suppose that, for each kind of object, we maintain some kind of table of associations between levels of illumination and corresponding values for luminance. Since we are unlikely to have an inborn knowledge of coal, snow, and the like, we must suppose that those tables are learned. As noted in Section 6.3.2, learning associations of that kind is part of what the SP system is designed to achieve.

Notice that any given table can only be applied if we have some idea of what kind of object we are looking at, otherwise we might see coal as if it was snow, or *vice versa*. There is some evidence that our perception of the lightness of an object does indeed depend on what we think the object is (Frisby and Stone, 2010, Chapter 16). In a similar way, our judgements of lightness seem to depend on our perceptions of how a given object is illuminated (Stone, 2012, Figure 1.10).

It seems likely that much of what has been said in this section about lightness constancy would also apply to colour constancy: the way we see the colour of an object to be fixed, despite wide variations in the colour of the incident light and corresponding variations in the colour of the light that is reflected from the object.

Since information compression is central in the SP theory, it is pertinent to mention that lightness constancy and colour constancy may each be seen as a means of encoding information economically. It is simpler to remember that a particular object is 'black' or 'red' than all the complexity of how its appearance changes in different lighting conditions.

7.4 The role of context in recognition

It is often remarked that we recognise things more easily in their familiar contexts than in unfamiliar ones, and this is confirmed in formal studies (see, for example, Bar and Ullman, 1993; Oliva and Torralba, 2007).

This observation makes sense in terms of the SP framework because any part of a multiple alignment may be a context for any other, and because of the way the system searches for a global optimum which embraces any given entity and its context. If, in our seascape example (Section 4.3), we see a beach and the sea then, in effect, we are primed to see boats—because, in that context, boats are likely to yield multiple alignments with better scores than, say, office furniture.

7.5 Ambiguity in perception

A less common observation is that, with some kinds of image, there is more than one plausible interpretation. An example is the 'young woman / old woman' picture of psychology text books.[27]

In the SP framework, this kind of ambiguity is accommodated in the way that, with some kinds of data, the system may create two or more multiple alignments that have good scores. An example in the area of natural language processing is the way the SP62 model can produce two parsings corresponding to both readings of the ambiguous sentence *Fruit flies like a banana*, as shown in Wolff (2006, Figure 5.1).[28]

7.6 Integration of vision with other senses and other aspects of intelligence

It is clear that in people and other animals, vision does not stand alone but works in close association with other senses. Our concept of a ship, for example, is an amalgam of images, sounds, smells, the flavour of food on board, textures of different surfaces, and so on.

In a similar way, vision works closely with other aspects of intelligence: different kinds of reasoning, learning, understanding and producing natural language, recalling information, and non-visual kinds of recognition.

Achieving these kinds of integration without undue complexity has been a central aim in the development of the theory. And in that development, many

[27] Another popular example is a picture that can be seen as either a duck or a rabbit.

[28] The given sentence is the second part of *Time flies like an arrow. Fruit flies like a banana.*, attributed to Groucho Marx.

candidate ideas have been rejected because they did not help to promote the simplification and integration of concepts.

To the extent that the theory achieves a combination of simplicity with versatility, it is down to three main things: representing all kinds of knowledge with 'patterns'; the multiple alignment concept as it has been developed in the SP theory; and the overarching role of information compression via the matching and unification of patterns.

8 Conclusion

Despite some limitations in how the SP theory is currently realised in computer models, it has what I believe are some useful things to say about several aspects of vision:

- Low level perceptual features such as edges or corners may be identified by the extraction of redundancy in uniform areas in a manner that is analogous to the run-length encoding technique for information compression, and comparable with the effect of lateral inhibition in the visual systems of animals.
- The concept of *multiple alignment* in the SP theory may be applied to the recognition of objects, and to scene analysis, with a hierarchy of parts and sub-parts, and at multiple levels of abstraction.
- The theory has potential for the unsupervised learning of visual objects and classes of objects, and suggests how coherent concepts may be derived from fragments. It provides an account of how we may discover objects via stereo matching and via motion.
- As in natural vision, both recognition and learning in the SP system is robust in the face of errors of omission, commission and substitution.
- The theory suggests how, via vision, we may piece together a knowledge of the three-dimensional structure of objects and of our environment that is good enough for many practical purposes, despite ambiguities in geometry.
- The theory provides an account of how we may see things that are not objectively present in an image, and how we may recognise something despite variations in the size of its retinal image.

- The theory has things to say about the phenomena of lightness constancy and colour constancy, about the role of context in recognition, and about ambiguities in visual perception.

A strength of the SP theory is that it is not simply a theory of vision. It provides for the integration of vision with other sensory modalities and with other aspects of intelligence such as reasoning, planning, and problem solving.

References

F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61:183–193, 1954.

M. Bar and S. Ullman. Spatial context in recognition. *Perception*, 25:324–352, 1993.

H. B. Barlow. Sensory mechanisms, the reduction of redundancy, and intelligence. In HMSO, editor, *The Mechanisation of Thought Processes*, pages 535–559. Her Majesty's Stationery Office, London, 1959.

H. B. Barlow. Trigger features, adaptation and economy of impulses. In K. N. Leibovic, editor, *Information Processes in the Nervous System*, pages 209–230. Springer, New York, 1969.

R. Brown. *A First Language: The Early Stages*. Penguin, Harmondsworth, 1973.

H. H. Bülthoff and S. Edelman. Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Acadamy of Sciences*, 89(1):60–64, 1992.

C. Farabet, C. Couprie, L. Najman, and Y. LeCun. Scene parsing with multiscale feature learning, purity trees, and optimal covers. In *Proceedings of the 29th International Conference on Machine Learning, Edinburgh, Scotland, UK, 2012*, 2012.

J. P. Frisby and J. V. Stone. *Seeing: The Computational Approach to Biological Vision*. The MIT Press, London, England, 2010.

K. S. Fu, editor. *Syntactic Pattern Recognition, Applications*. Springer-Verlag, Berlin, 1977.

W. L. Gehringer and E. Engel. Effect of ecological viewing conditions on the ames' distorted room illusion. *Journal of Experimental Psychology: Human Perception and Performance*, 12(2):181–185, 1986.

A. Glennerster, S. J. Gilson, L. Tcheang, and A. J. Parker. Perception of size in a 'dynamic Ames room'. *Journal of Experimental Psychology: Human Perception and Performance*, 3(9):490a, 2003. doi: 10.1167/3.9.490.

M. Gold. Language identification in the limit. *Information and Control*, 10: 447–474, 1967.

F. Han. Bottom-up/top-down image parsing by attribute graph grammar. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV 2005), 17-21 Oct. 2005*, volume 2, pages 1778–1785, 2005.

B. Julesz. *Foundations of Cyclopean Perception*. Chicago University Press, Chicago, 1971.

E. H. Lenneberg. Understanding language without the ability to speak. *Journal of Abnormal and Social Psychology*, 65:419–425, 1962.

M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 2009.

D. Marr. *Vision: a Computational Investigation into the Human Representation and Processing of Visual Information*. The MIT Press, London, England, 2010. This book was originally published in 1982 by W. H. Freeman and Company.

D. Marr and T. Poggio. A computational theory of human stereo vision. *Proceedings of the Royal Society of London. Series B*, 204(1156):301–328, 1979.

A. Oliva and A. Torralba. The role of context in object recognition. *Trends in cognitive sciences*, 11(12):520–527, 2007.

F. Ratliff and H. K. Hartline. The response of limulus optic nerve fibres to patterns of illumination on the receptor mosaic. *Journal of General Physiology*, 42:1241–1255, 1959.

Q. Y. Shi. Parsing and translation of (attributed) expansive graph languages for scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5(5):472–485, 1983.

R. J. Solomonoff. A formal theory of inductive inference. Parts I and II. *Information and Control*, 7:1–22 and 224–254, 1964.

J. V. Stone. *Vision and Brain: How We Perceive the World*. The MIT Press, London, England, 2012.

R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer, London, 2011. ISBN 978-1-84882-934-3.

M. J. Tarr. Rotating objects to recognize them: A case study of the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bulletin and Review*, 2(1):55–82, 1995.

M. J. Tarr and S. Pinker. Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, 21(2):233–282, 1989.

G. von Békésy. *Sensory Inhibition*. Princeton University Press, Princeton, NJ, 1967.

J. G. Wolff. The discovery of segments in natural language. *British Journal of Psychology*, 68:97–106, 1977. See: bit.ly/Yg3qQb.

J. G. Wolff. Language acquisition and the discovery of phrase structure. *Language & Speech*, 23:255–269, 1980. See: `www.cognitionresearch.org/lang_learn.html#wolff_1980`.

J. G. Wolff. Language acquisition, data compression and generalization. *Language & Communication*, 2:57–89, 1982. See: `www.cognitionresearch.org/lang_learn.html#wolff_1982`.

J. G. Wolff. Learning syntax and meanings through optimization and distributional analysis. In Y. Levy, I. M. Schlesinger, and M. D. S. Braine, editors, *Categories and Processes in Language Acquisition*, pages 179–215. Lawrence Erlbaum, Hillsdale, NJ, 1988. See: bit.ly/ZIGjyc.

J. G. Wolff. *Unifying Computing and Cognition: the SP Theory and Its Applications*. CognitionResearch.org, Menai Bridge, 2006. ISBNs: 0-9550726-0-3 (ebook edition), 0-9550726-1-1 (print edition). Distributors, including Amazon.com, are detailed on bit.ly/WmB1rs. The publisher and its website was previously CognitionResearch.org.uk.

J. G. Wolff. Towards an intelligent database system founded on the SP theory of computing and cognition. *Data & Knowledge Engineering*, 60:596–624, 2007. See: bit.ly/Yg2onp.